

Analyses of Cardinal Auctions

Mangesh Gupte*
Google Inc
mangesh@cs.rutgers.edu

Darja Krushevskaja
Rutgers University
darja@cs.rutgers.edu

S. Muthukrishnan
Rutgers University
muthu@cs.rutgers.edu

March 3, 2013

Abstract

We study *cardinal auctions* for selling multiple copies of a good, in which bidders specify not only their bid or how much they are ready to pay for the good, but also a cardinality constraint on the number of copies that will be sold via the auction. We perform first known *Price of Anarchy* type analyses with detailed comparison of the classical Vickrey-Clarke-Groves (VCG) auction and one based on minimum pay property (MPP) which is similar to Generalized Second Price auction commonly used in sponsored search. Without cardinality constraints, MPP has the same efficiency (total value to bidders) and at least as much revenue (total income to the auctioneer) as VCG; this also holds for certain other generalizations of MPP (e.g., prefix constrained auctions, as we show here). In contrast, our main results are that, with cardinality constraints, (a) equilibrium efficiency of MPP is $1/2$ of that of VCG and this factor is tight, and (b) in equilibrium MPP may collect as little as $1/2$ the revenue of VCG. These aspects arise because in presence of cardinality constraints, more strategies are available to bidders in MPP, including bidding above their value, and this makes analyses nontrivial.

1 Introduction

Assume we have n bidders and m identical items to sell via an auction, s.t. $n \geq m$. What auction is suitable? In particular, there are three decisions to be made:

- how many items to sell: k^* ,
- how to allocate k^* items: $a(\cdot)$ and
- how to price each of them: $p(\cdot)$.

We consider the case of *negative externality* when the number of bidders who win and are allotted the item affects the value of the item to each of the winners. In particular, each bidder is not interested in a copy if the number of copies eventually allocated exceeds her threshold. *Cardinal auctions* explicitly incorporate this externality into the bidding language. In this paper we study efficiency and revenue tradeoffs for such auctions.

2 Model

In cardinal auctions, there are n buyers competing for at most $m \leq n$ identical copies of an item in the auction. Each buyer wants to buy exactly one copy and has two *private* numbers v_i and

*This work was done while the author was graduate student at Rutgers University.

k_i . Auctioneer has no prior information about values of buyers. The utility $u_i(v_i, k_i)$ that bidder i derives from the auction is

$$u_i(v_i, k_i) = \begin{cases} x_i v_i - p_i & \text{if number of copies sold is less than } k_i \\ -\infty & \text{otherwise} \end{cases}$$

where $x_i \in \{0, 1\}$ is indicator variable that shows whether i was allotted a copy or not, and p_i is the price at which i obtains it.

Buyers express their preferences through *2 dimensional bid* (b_i, l_i) where b_i is the maximum amount buyer i is willing to pay if at most l_i copies are allocated. Note, that b_i may differ from v_i , and l_i from k_i . Once the auctioneer gathers all the bids she has to decide on optimal number of copies k^* , allocation of k^* copies according to function $a(\cdot)$ and payments according to pricing function $p(\cdot)$. In mechanisms we will consider, no bidder i will be a winner if $l_i < k^*$.

Motivating Scenarios. An important motivation arises in auctions for online advertisements (ads). Consider *display ads*, or visual ads, on a webpage. Advertisers whose ad is shown on the page compete for attention of the viewers. Clearly, the *number* of ads shown is an important feature, e.g., publishers recognize that showing fewer ads helps¹. Currently, this cardinality is largely determined by the publisher of the web page, who may choose to make it exclusive showing only one ad, but in many cases mixes several. They choose the number of ads on a page based on variety of techniques from machine learning to user studies, esthetics of UI design and revenue maximization. This approach does not let advertisers influence how many ads appear with their own; hence, they bid depending on the average of their values over the possible number of ads that might appear on that page. This induces inefficiencies and potential revenue loss. Cardinal auctions are an alternative. They let advertisers explicitly specify how many other advertisers may appear with their ad on a given page.

Cardinal auctions are also suitable in a variety of other instances:

- Say we can produce a collectors item such as a signed copy of an album or a book. The more exclusive the copy is, the more valuable it is to the possessor. How many copies shall we produce? While traditionally this is determined by estimating the demand function, one can imagine an auction-based method, where bidders can specify in some way the value of the item to them as a function of how many copies are made and sold.
- Consider a situation that arises in a data exchange such as BlueKai² where certain pertinent data about a user is sold for ads targeting. The data may be sold to any number of advertisers for targeting, but in some cases, the more the information is shared, the less value it gives to the advertisers. Hence, when data is sold via auction, advertisers may wish to be able to influence how many of others get access to the data.

2.1 Auctions

Allocation \mathcal{A} is the set of k^* winners who obtain a copy. We consider set of *feasible* allocations: allocation \mathcal{A} is *feasible* if $\{l_i \geq |\mathcal{A}|, k_i \geq |\mathcal{A}| : \forall i \in \mathcal{A}\}$. The total *efficiency* $\mathcal{E}_{\mathcal{A}}$ of allocation \mathcal{A} is the sum of values of allotted bidders or $\sum_{i \in \mathcal{A}} v_i$.

There are two natural auctions to consider.

VCG_{CA}. VCG_{CA} is a straightforward extension of the standard Vickrey-Clarke-Groves (VCG) auction [Gro71, Vic61, Cla73]. VCG_{CA} is *truthful*, i.e., bidder's bid their true valuations. This

¹<http://www.technologyreview.com/web/25827/?a=f>

²<http://www.bluekai.com/>

is a well-known property of VCG mechanism. Thus we can assume that bidders submit their bids as (v_i, k_i) . VCG_{CA} chooses the feasible allocation that maximizes the total sum of values: $\mathcal{E}^* = \max_{\mathcal{A}} \sum_{i \in \mathcal{A}} v_i$. Let \mathcal{E}_{-i} be highest efficiency achievable without bidder i , then the price for bidder i is

$$p_i^{VCG} = \mathcal{E}_{-i} - \mathcal{E}^* + v_i \quad (1)$$

VCG_{CA} , as generally known, can have low revenue. Furthermore, with cardinality constraints its outcomes are not *envy-free*: losing bidder would agree to purchase a copy for the price higher than what is being asked from winners. For more intuition, consider the following example:

Example 2.1. *There are 3 bidders with true valuations: $A = (100, 1)$, $B = (90, 2)$ and $C = (80, 2)$. For this setting VCG_{CA} will identify winning allocation with bidders B and C in it since the total efficiency is $90 + 80 = 170 > 100$). It will charge B amount $p_B = 100 - 170 + 90 = 20$ and C amount $p_C = 100 - 170 + 80 = 10$. Thus total collected payment is 30. However, bidder A will envy this low payment of 30.* \square

MPP_{CA}. MPP_{CA} was introduced in [Mut09] and is based on the *minimum pay property* for the outcome: auction requires every buyer to pay no more than what she would have bid, if she knew all other bids, to get the exact same assignment she got. To calculate prices, let winning allocation be \mathcal{A}^* and \mathcal{A}_2 the allocation that gives second highest sum of bids after \mathcal{A}^* . Let winning bids in \mathcal{A}^* be sorted top-down in decreasing order of bids, $b_1 \geq b_2 \geq \dots$. The i th winner pays price

$$p_i^{MPP} = \max\left\{\sum_{j \in \mathcal{A}_2} b_j - \sum_{j \in \mathcal{A}^*} b_j + b_i, b_{i+1}\right\}$$

The price consists of two components. The first term is the minimum amount i needs to bid to ensure that the allocation \mathcal{A}^* is the winner, and the second term is the minimum bid to get above the $i + 1$ 'st largest bid. The overall price is the maximum over both.

MPP auction is inspired by *Generalized Second Price* (GSP) auction used by many popular search engines [EOS05, Gro71, AGM06] to determine placement of advertisements (ads) on the page. In *GSP* there are n advertisers bidding for m advertisement slots. Each slot i has associated *click through rate* (CTR) with it, or probability of being clicked, denoted by $\alpha_i \in (0, 1)$. Slots are ordered in decreasing order of CTR's: $\alpha_i > \alpha_j$ for $i < j$. Advertiser i has private valuation v_i , which expresses the value of getting a click. To participate in the auction advertiser submits bid b_i that indicates maximum payment she is willing to make. Auctioneer receives all bids, and assigns advertisers to slots in decreasing order of their bids. For convenience, let us renumber advertisers in decreasing order of their bids, then, advertiser i is assigned to slot i with CTR α_i . Payment of advertiser i is $p_i = \frac{\alpha_{i+1}}{\alpha_i} b_{i+1}$, and is charged only if the ad is clicked.

MPP_{CA} naturally generalizes *GSP* auction and in absence of cardinality constraints is the special case of *GSP* without click through rates, or when $\alpha_i = 1$ for all i .

2.1.1 Analysis of Auctions

Unlike VCG_{CA} , MPP_{CA} is not truthful. While bidder cannot benefit from misreporting her cardinality preference³, she can improve her utility by reporting $b_i \neq v_i$. Consider the following example:

Example 2.2. *Consider 3 bidders with their true valuations: $A = (100, 1)$, $B = (80, 2)$ and $C = (70, 2)$. Auctioneer runs MPP_{CA} . If auctioneer receives truthful bids, then she will choose*

³we will show this in a separate argument later on

allocation of 2 bidders: (B, C) , and charge them 70 and 20, respectively. Utility of bidder B is $u_B = 80 - 70 = 10$. However, bidder B can improve it by lowering her bid to $(40, 2)$. Then, allocation is the same (bidders B and C), but payments are different: payment of B is 30, and payment of C is 60. Now, utility of B is $u'_B = 80 - 30 = 50 > 10$. Hence, bidder B benefits from submitting untruthful bid $b_B < v_B$. \square

MPP_{CA} can have many outcomes, we consider only bid vectors that are in *Nash equilibria*, that is, for every bidder $i \in \mathcal{A}$ the following inequalities hold:

$$\begin{aligned} v_i - p_{i,\mathcal{A}} &\geq v_i - p_{j,\mathcal{A}'} & \forall \mathcal{A}' \in \mathcal{F}, \mathcal{A} \neq \mathcal{A}' \\ v_i - p_{i,\mathcal{A}} &\geq v_i - p_{j,\mathcal{A}} & \forall j \neq i \end{aligned}$$

There is a set of Nash equilibria efficiencies of $\Sigma_1, \Sigma_2, \dots$. Define $\mathcal{E}_{\min} = \min_i \Sigma_i$ and $\mathcal{E}_{\max} = \max_i \Sigma_i$. Then, *price of anarchy* is defined as $PoA(\cdot) = \frac{\mathcal{E}_{\max}}{\mathcal{E}_{\min}}$. Observe, that $\mathcal{E}_{\max} = \mathcal{E}^*$. Thus, $PoA(VCG_{CA}) = 1$, furthermore in order to evaluate performance of MPP_{CA} , one can use VCG_{CA} as a benchmark, and compare efficiency of VCG_{CA} outcome with that of the worst outcome of MPP_{CA} :

$$PoA(MPP_{CA}) = \frac{\mathcal{E}^*}{\mathcal{E}_{\min}}.$$

To analyze revenue we use VCG_{CA} as the benchmark for consistency, and compare it to the revenue collected by MPP_{CA} .

Bidders are strategic and their goal is to maximize their utility. Their behavior, or *strategy*, is determined by mechanism. Strategy s is *weakly dominant* if regardless of what other bidders do, strategy s gets a player utility that is at least as high as utility obtained by playing any other strategy. Strategy s is (*strictly*) *dominant* if utility of playing strategy s is strictly larger than playing any other strategy, regardless what other bidders do. We consider two types of bidders:

- *Conservative bidders* do not bid over their value, i.e. $b_i \leq v_i$, hence they do not risk paying more than their true valuation and getting negative utility.
- *Rational bidders* can bid above their true valuation v_i in equilibria iff the payment p_i does not exceed v_i . Equilibria that contain such bids are fragile, because bidder can get negative utility if some other bidder changes her bid. Such equilibria help us explore the properties of possible outcomes.

2.2 Our Results

We perform first known analyses of cardinal auctions, and compare efficiency and revenue of VCG_{CA} vs MPP_{CA} in equilibrium.

- (*Efficiency*) We show that PoA is 1 for conservative bidders. For rational bidders, we show that PoA is 2 and this is optimal.

As noted before, without cardinal constraints l_i , MPP_{CA} becomes GSP without click through rates. In that case, it is a dominant strategy for bidders to be conservative, and PoA is 1. Further, we show that, even with a slightly different bidding language of *prefix constraints* (defined later precisely), it is still (weakly) dominant strategy for bidders to be conservative, and PoA of prefix auctions is still 1. It is interesting that with cardinality constraints, there is provable loss of efficiency in equilibrium.

- (*Revenue*) We show that for conservative bidders revenue of MPP_{CA} is always at least that of VCG_{CA} , and for rational bidders, revenue of MPP_{CA} may be only 1/2 of that of VCG_{CA} .

In contrast, without cardinality constraints, MPP_{CA} has larger revenue than VCG_{CA} , no matter the nature of bidders.

In both analyses of efficiency and revenue, the central technical challenge is that MPP_{CA} pricing has two components, first ensures that the eventual allocation is “better” than others with fewer, or more items, and the other is the impact of the bidder in the chosen allocation. The role of position component of price of MPP_{CA} has been studied extensively in analyses of GSP [LT10, LPL11]. The cardinality component of the pricing of MPP_{CA} allows bidders to bid over their true valuation, this induces the nontrivial PoA properties, complicates the analyses, and is novel.

2.3 Related Work

The work closest to ours is [GS10, JS11, Mut09]. In [Mut09] authors introduce auction with cardinal externality, formulate and motivate pricing, and give efficient algorithms for calculating the allocation and pricing. However, they do not perform PoA type analysis we do here.

[GS10, JS11] introduce auction with negative externality: valuation of bidder i is a pair (v_i^E, v_i^M) , where v_i^E is value of the bidder for being *exclusively* allocated, and v_i^M is valuation if bidder i for being allocated among with other bidders. The auction determines allocation type (exclusive vs non exclusive), set of bidders allotted, and prices at which they get the item. The authors consider various pricing schemes including variations of VCG and MPP and provide analysis of efficiency and revenue in equilibrium. Our work differs from [GS10, JS11] in the bidding language. The bidding language in [GS10, JS11] can not specify preference for the number of copies being sold, while our bidding language does not allow to specify more than a single bid. Thus, the bidding languages are incomparable. We believe both languages are natural and interesting.

Another relevant auction was presented in [AFM07], where authors consider the auction for the ordered set of items, in which bidders specify the largest prefix of allocation in which they want to participate: valuation of the bidder is (v_i, k_i) , where bidder i has valuation v_i only if she is in top k_i allotted bidders. In their model bidders have no influence on valuations of each-other or the size of final allocation. Authors present two auctions: one based on VCG and the other based on MPP. The paper shows that MPP based auction can achieve efficiency of VCG, however do not present PoA type results. In this paper we present analysis of efficiency and revenue for this model and contrast it with our main results of efficiency and revenue for cardinal auctions.

In general, there are limited PoA style analyses of auctions. Closely related to cardinal auctions are sponsored search auctions where each item has an associated click through rate. A recent paper performs PoA style analysis in this setting: in [LT10] authors show that it is weakly dominant strategy for bidders to bid conservatively and PoA of $\frac{1+\sqrt{5}}{2}$ for Nash equilibria of GSP with conservative bidders.

There is extensive body of literature on *multi-unit auctions*: auctioneer has m items to sell among n buyers. The problem was studied at least since [SS71], however much of that work does not accommodate externalities, and considers the number of items to sell k as an input to the model. A more recent phenomena is auctions with unlimited supply, or digital good auctions [GH03, GHK⁺06]. Here, auctioneer has unlimited supply of the good to sell, she needs to find optimal price p^* that determines both allocation and uniform price. However buyer valuation v_i ’s are independent of the outcome of the allocation, unlike what we study here.

Externalities have been studied extensively in Economics and Computer Science. [HIMM11] considers model where buyers experience positive externalities once sufficient number of their friends

are also allocated the copy of the item. In [JMS96] authors consider situation when winning buyer subjects other participants to negative externality. [AC08] studies the problem of allocating a pair of goods among group of sellers who have pairwise externalities to each other.

[PKK11] considers auction for sharable goods. In their setting, valuation of the bidder is multidimensional and represents value of strict ownership, sharing value, and that of no allocation. They investigate unique perfect Bayesian equilibrium. All of these works consider externality as an input to the problem or model value of a bidder as some function of number of items allocated. All bidders are treated equally and no bidder has control over allocation in which they participate. This is different from cardinal auctions we study here where the bidding language explicitly allows bidders to specify the cardinal constraints.

At the highest level, it will be of interest to study auctions for the model in which buyer can explicitly define her valuation for each possible allocation size and the set of different identities of co-winners. However, such exponential input can not be processed or analyzed efficiently, and even in restricted cases, at least as hard as various multidimensional auctions.

In practical terms, within the context of advertising, there were machine learning approaches to estimating the optimal number of ad slots on the page. In [SKN⁺10] authors consider several utility functions that incorporate user experience in order learn number of slots that optimizes utility function. [BCF⁺08] formulates a binary problem: given a set of relevant ads the goal is to decide whether system benefits from showing ads or should it not show the ads and by that benefit in a long run. These are interesting directions, different from our approach which relies on the bidders to address the issue (at least in the short run).

3 Preliminary Observations

Our first observation is regarding truthfulness of the cardinal constraint.

Lemma 3.1. *In MPP_{CA} , bidders truthfully reveal their private k_i 's, that is $l_i = k_i$ for all i in Nash equilibria of MPP_{CA} auction.*

Proof. Consider any bidder i who participates in winning allocation of size k^* . For contradiction let their revealed $l_i > k_i$. There are two possibilities: (1) $k^* \leq k_i$, then bidder i would have identical utility by reporting k_i instead of l_i ; (2) $k^* > k_i$, bidder now has utility $-\infty$ which is worse than utility from reporting k_i truthfully. Hence, bidder does not have incentive to submit $l_i > k_i$.

Consider same bidder i , and say for contradiction that their revealed $l_i < k_i$. Again, there are two cases: (1) $k^* \leq l_i$, then bidder i would have identical utility by reporting k_i instead of l_i ; (2) with some positive probability $k_i \geq k^* > l_i$, bidder now has utility 0 which is worse than utility from reporting k_i truthfully. Hence, bidder does not have incentive to submit $l_i < k_i$. \square

Our second observation concerns bidding behavior of *losing bidders*. There exist bids which are in equilibrium, but the efficiency is bounded away from the maximum achievable efficiency by an arbitrarily large factor. Consider the case of three bidders $(100, 1), (75, 2), (75, 2)$, who make the following bids: $(100, 1), (1, 2), (1, 2)$. This set of bids forms Nash equilibrium since no bidder by herself has an incentive to change her bid. However, the efficiency of the resulting allocation is 100, compared to the optimum efficiency of 150. This is due that losing bidders can arbitrarily shade their bids. This is a common problem, previously faced in [GS10]. Like [GS10], we will henceforth assume that losing bidders bid their true valuation.

4 Efficiency Analyses

Theorem 4.1. *With conservative bidders, MPP_{CA} 's allocation has the same total value as VCG_{CA} in Nash equilibrium.*

Proof. Let $\sigma = \sigma_{VCG}$ be the set of winning bidders that maximizes efficiency and μ be the set of winning bidders under MPP_{CA} in Nash equilibria. Let $\{b_i | i \in \mu\}$ be a set of equilibrium bids under MPP_{CA} . Since MPP_{CA} chooses the set of bidders who maximizes total sum of bids, then it must be true that $\sum_{j \in \mu} b_j \geq \sum_{j \in \sigma} b_j$. Since σ is feasible, $\sum_{j \in \mu} b_j \geq \sum_{j \in \sigma} b_j$. Hence

$$\begin{aligned}
& \sum_{i \in \mu \setminus \sigma} b_i + \sum_{i \in \mu \cap \sigma} b_i \geq \sum_{j \in \sigma \setminus \mu} b_j + \sum_{j \in \mu \cap \sigma} b_j \\
\Rightarrow & \sum_{i \in \mu \setminus \sigma} b_i \geq \sum_{j \in \sigma \setminus \mu} b_j \\
\Rightarrow & \sum_{i \in \mu \setminus \sigma} v_i \geq \sum_{j \in \sigma \setminus \mu} v_j \\
\Rightarrow & \sum_{i \in \mu} v_i \geq \sum_{j \in \sigma} v_j
\end{aligned} \tag{2}$$

In (2) we use the assumption on the right hand side that losers bid at least their true valuations and on the left hand side that bidders are conservative. Then the only possibility is that total value of μ equals that of σ . \square

It follows that the PoA of MPP_{CA} is 1 for conservative bidders.

Theorem 4.2. *For rational bidders, PoA of MPP_{CA} is 2 and this is tight.*

Proof. Let μ and σ denote the set of bidders chosen by the allocation of MPP_{CA} and VCG_{CA} respectively. Let μ_2 denote set of bidders who belong to second best allocation of MPP_{CA} .

If $\sigma = \mu$, then the efficiency is 1 and we are done. Otherwise,

$$\begin{aligned}
& \mu = \sum_{i \in \mu} b_i > \sum_{j \in \sigma} b_j \\
\Rightarrow & \sum_{i \in \mu \setminus \sigma} b_i + \sum_{i \in \mu \cap \sigma} b_i > \sum_{j \in \sigma \setminus \mu} b_j + \sum_{i \in \mu \cap \sigma} b_i \\
\Rightarrow & \sum_{i \in \mu \setminus \sigma} b_i > \sum_{j \in \sigma \setminus \mu} b_j \\
\Rightarrow & \sum_{i \in \mu \setminus \sigma} b_i > \sum_{j \in \sigma \setminus \mu} v_j
\end{aligned} \tag{3}$$

where to get (3) we use assumption that losing bidders bid at least their value. Remainder of the proof deviates from the conservative bidder case as we can not bound the left hand side for rational like we did with conservative bidders.

Without loss of generality, assume that the bidders in $\mu \setminus \sigma = \{b_1, b_2, \dots, b_k\}$ are ordered in non-increasing order of bids, i.e., $b_1 \geq b_2 \geq \dots \geq b_k$. To lowerbound the payment of the highest bidder, we start by working on one of the components of the pricing:

$$\begin{aligned}
\sum_{i \in \mu_2} b_i - \sum_{i \in \mu} b_i + b_1 & \geq \sum_{i \in \sigma} b_i - \sum_{i \in \mu} b_i + b_1 \\
& \geq \sum_{i \in \sigma \setminus \mu} v_i - \sum_{2 \leq i \leq k} b_i + b_1
\end{aligned} \tag{4}$$

where we get the first term of (4) from Eq. 3. Notice, that if highest bidder i belongs to σ , then there is at least one bidder in μ who pays more than her value and gets negative utility. Hence, for allocation to be in Nash equilibria highest bidder i must belong to $\mu \setminus \sigma$, and we can exclude highest bidder from second term of (4). Hence,

$$p(b_1) \geq \max \left\{ b_2, \sum_{i \in \sigma \setminus \mu} v_i - \sum_{2 \leq i \leq k} b_i \right\}$$

For other bidders, we bound MPP_{CA} payment by the $p(b_i) \geq b_{i+1}$. Using these, we get a lower bound on the total revenue of MPP_{CA} as follows:

$$\begin{aligned} \sum_{1 \leq i \leq k} p(b_i) &= p(b_1) + \sum_{2 \leq i \leq k} p(b_i) \\ &\geq \max \left\{ b_2, \sum_{i \in \sigma \setminus \mu} v_i - \sum_{2 \leq i \leq k} b_i \right\} + \sum_{2 \leq i \leq k} b_{i+1} \\ &\geq \max \left\{ b_2, \sum_{i \in \sigma \setminus \mu} v_i - \sum_{2 \leq i \leq k} b_i + \sum_{3 \leq i \leq k} b_i \right\} \\ &= \max \left\{ b_2, \sum_{i \in \sigma \setminus \mu} v_i - b_2 \right\} \geq \frac{1}{2} \sum_{i \in \sigma \setminus \mu} v_i \end{aligned}$$

Since, the bidders are rational,

$$\sum_{i \in \mu \setminus \sigma} v_i \geq \sum_{i \in \mu \setminus \sigma} p_i = \sum_{1 \leq i \leq k} p(b_i)$$

and chaining with the previous equation, we get that $\sum_{i \in \mu \setminus \sigma} v_i \geq \frac{1}{2} \sum_{v_i \in \sigma \setminus \mu} v_i$.

Tightness. To see that the bound is tight consider three bidders with the following valuations: $(100, 1)$, $(50, 2)$, $(\epsilon, 2)$ and the bids they place are $(100, 1)$, $(100, 2)$, $(50, 2)$. Bids form NE, and none of the bidders can improve her utility acting on her own. MPP_{CA} will allocate bidders $(2, 3)$ and charge them 50 and 0 respectively. PoA is then $\frac{50+\epsilon}{100} \approx \frac{1}{2}$. \square

Contrasts with other auctions As mentioned earlier, without cardinality constraints MPP_{CA} becomes GSP auction without click-through rates. It is known, that in that case PoA of GSP is 1. To highlight our result further, we consider PoA of *prefix auctions* [AFM07] and show that it is also 1.

The model is as follows. There are n ordered identical items to sell. There are m bidders, each bidder i has two private values v_i and k_i . Utility u_i of bidder i is $v_i - p_i$ if she obtains *any* of the first k_i copies, and it is $-\infty$ otherwise. Notice that now, bidder i has positive utility even if more than k_i copies are auctioned. In the auction, each bidder i submits a pair (b_i, l_i) : b_i is the maximum they are willing to pay, if they are allotted one of the first l_i copies.

We consider two different auctions: $pVCG$ and $pGSP$:

pVCG. $pVCG$ is extension of VCG and is truthful. In the auction bidders submit their true valuations (v_i, k_i) to auctioneer, upon receipt of bids auctioneer creates feasible allocation that maximizes total efficiency and calculates payments using Eq. 1.

pGSP. $pGSP$ is iterative second price (SP) auction from first copy of the item to the last, where for each item, we run a SP auction among bidders who are not yet assigned a copy but who still have nonnegative utility from obtaining one.

Notice, that bidder cannot benefit by submitting bid $b_i > v_i$, hence it is weakly dominant strategy for bidders to bid *conservatively*. Similarly to MPP_{CA} it is dominant strategy for the bidder to report her preference k_i truthfully. The argument is identical to Lemma 3.1.

Unlike second price auction $pGSP$ is not truthful. Thus, similarly to MPP_{CA} , we analyze bid vectors that are in Nash equilibria. For $pGSP$, Nash equilibrium is defined as follows. For each i ,

$$v_i - p_i \geq v_i - p_j \quad \forall j \neq i \text{ and } j \leq \min\{k, k_i\}$$

Here we give PoA results for prefix auctions without click-through rates.

Theorem 4.3. *PoA of $pGSP$ is 1.*

Proof. Let σ be set of bidders allocated by $pVCG$ and δ be set of bidders allocated by $pGSP$ in Nash equilibria. Then, similarly to Eq. 3 we have

$$\sum_{i \in \sigma \setminus \delta} v_i \geq \sum_{j \in \delta \setminus \sigma} v_j$$

It is possible only if

$$\exists i \in \sigma \setminus \delta \quad \forall j \in \delta \setminus \sigma \text{ s.t. } v_i > v_j \quad (5)$$

Assume, it is not the case. Then, $\exists j \in \delta \setminus \sigma \quad \forall i \in \sigma \setminus \delta \text{ s.t. } v_j > v_i$. However, this is possible only if bidder v_j cannot replace any of bidders $i \in \sigma \setminus \delta$, otherwise v_j would improve \mathcal{E}^* . This, in turn, is possible if and only if $|\{i | i \in \sigma \setminus \delta\}| = 0$ that would imply that $|\sigma| < |\delta|$. However, it gives a contradiction, as efficiency of \mathcal{E}^* could be improved by adding v_j to it.

If (5) is true, then there must be a losing bidder l who can raise her bid, enter the allocation and as the result improve her utility. That gives a contradiction. Thus, total values of σ and δ are identical. \square

5 Revenue Analyses

Let $Rev(X)$ be revenue generated by mechanism X . We show two results.

Theorem 5.1. *With conservative bidders, $Rev(MPP_{CA}) \geq Rev(VCG_{CA})$.*

Proof. Let σ be the allocation with maximum total value (hence, the value attained by VCG_{CA}), μ be the allocation of MPP_{CA} in equilibrium, μ_2 be the set of bidders who participate in second best allocation of MPP_{CA} and σ_{-i} the set of bidders that gives the largest total value allocation when bidder i is not present.

Consider payments of each bidder $i \in \sigma \cap \mu$ under VCG_{CA} and MPP_{CA} :

$$p_i^{VCG} = \sum_{j \in \sigma_{-i}} v_j - \sum_{j \in \sigma} v_j + v_i$$

$$p_i^{MPP} = \max\{b_{i+1}, \sum_{j \in \mu_2} b_j - \sum_{j \in \mu} b_j + b_i\} \geq \sum_{j \in \mu_2} b_j - \sum_{j \in \mu} b_j + b_i$$

Since bidders are conservative and losers bid their values,

$$\begin{aligned}
p_i^{VCG} &= \sum_{j \in \sigma - i} v_j - \sum_{j \in \sigma} v_j + v_i \leq \sum_{j \in \mu_2} v_j - \sum_{j \in \mu} b_j + b_i \\
&= \sum_{j \in \mu_2 \setminus \mu} b_j + \sum_{j \in \mu_2 \cap \mu} b_j - \sum_{j \in \mu \cap \mu_2} b_j - \sum_{j \in \mu \setminus \mu_2} b_j + b_i = p_i^{MPP}
\end{aligned}$$

Now, consider payments of all such bidders $i \in \sigma \setminus \mu$ or $i \in \mu \setminus \sigma$. This is possible, when there are 2 allocations of different size that have equally high efficiency, let's denote them by \mathcal{A}_σ and \mathcal{A}_μ . If bidder i is present in only one of allocations, then her payment is $p_i^{MPP} = p_i^{VCG} = v_i$. Payment of VCG_{CA} follows from definition. Observe, that bidder i submits truthful bid in MPP_{CA} , because otherwise she will be not be in winning configuration. Now, one can derive the payment from definition. \square

With rational bidders, we show that revenue of MPP_{CA} can be as low as half of that of VCG_{CA} .

Example 5.1. Consider 3 bidders with the following valuations $A = (100 + \epsilon, 1)$, $B = (50, 2)$ and $C = (50, 2)$. Rational bidders can converge to bids $A = (100, 1)$, $B = (100, 2)$ and $C = (50, 2)$ respectively. VCG_{CA} gets truthful bids and chooses allocation consisting of bidder A and her payment is 100, while MPP_{CA} chooses allocation with bidders $(B$ and $C)$ and prices them 50 and 0 respectively, achieving exactly half of revenue of VCG_{CA} . \square

As in case with efficiency, revenue of cardinal auctions is also surprising in contrast with other auctions. It is believed, that one of the reasons to use MPP auctions is to improve revenue, e.g., revenue of GSP without click-through rates is always at least as much as that of VCG . This is also true for modification presented in [GS10]. Likewise, for prefix auctions, this continues to hold.

Theorem 5.2. In equilibrium, $Rev(pVCG) \leq Rev(pGSP)$.

Proof. Let \mathcal{A} be the allocation of $pGSP$ (or $pVCG$). Consider payment of bidder $i \in \mathcal{A}$. Let l be the bidder who enters allocation \mathcal{A} if i leaves it. If no such bidder exist, let l be a bidder with valuation $v_l = 0$ and $k_l = n$. Then payment of bidder i in $pVCG$ is $p_i^{pVCG} = \mathcal{E}_{-i} + \mathcal{E} + v_i = v_l$ and payment of bidder i in $pGSP$ allocation $p_i^{pGSP} = \max\{b_{i+1}, b_l\}$.

Payment is minimized when $p_i^{pGSP} = b_l$. $b_l = v_l$, since bidders are conservative, and l is losing bidder. Hence, $Rev(pVCG) \leq Rev(pGSP)$. \square

In contrast to $pVCG$ and VCG that have lower revenue than the corresponding versions of MPP, for cardinal auctions, we have shown that in some cases VCG_{CA} may have more revenue than MPP_{CA} . This is because cardinal constraints enable richer strategies for bidders in particular strategy of rational bidders who can bid above their value.

6 Concluding Remarks and Future Directions

We consider the problem of selling identical copies of an item via an auction in which the number of copies sold is unknown *a priori*, and valuation of a bidder depends on the total number of winners. This scenario is motivated by number of ads on a page or number of parties that get access to certain information. While there are many ways to solve this problem, we consider *cardinal auctions* in which the bidding language lets buyers explicitly bid on the maximum number of winners allowed. Our work analyzes cardinal auctions of MPP_{CA} and VCG_{CA} for revenue and efficiency tradeoffs

in equilibria, and shows that they are quite different from the case without the cardinal externality. We find that MPP_{CA} which is inspired by widely used Generalized Second Price (GSP) auction has surprising properties. In case of rational bidders efficiency of MPP_{CA} is half of that of VCG_{CA} . At the same time, in the worst case MPP_{CA} can collect only half of revenue of VCG_{CA} .

There are many open directions to pursue. For example, in display ads, slots may differ in terms of their location and dimensions, as well as click through rates. We need to extend the study of cardinal auctions to auctions for configurations of display ads with varying quality scores or with varying click through models.

Externality is a richer phenomenon than we have studied here. For instance, the value for a bidder might depend not only on the number of other possessors, but also on their identity, quality, etc. Further, one can consider bidding languages which go beyond the step function we have adopted here, for example, by letting bidders specify their value for each potential number of winners. Studying such notions of externalities and bidding languages is an active area in Economics and problems are still open.

From a technical point of view, we would like to extend our analysis to Bayesian case and study dynamics of cardinal auctions.

References

- [AC08] J. Aseff and H. Chade. An optimal auction with identity-dependent externalities. *The RAND Journal of Economics*, 39(3):731–746, 2008.
- [AFM07] G. Aggarwal, J. Feldman, and S. Muthukrishnan. Bidding to the top: Vcg and equilibria of position-based auctions. *Approximation and Online Algorithms*, pages 15–28, 2007.
- [AGM06] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 1–7. ACM, 2006.
- [BCF⁺08] A. Broder, M. Ciaramita, M. Fontoura, E. Gabrilovich, V. Josifovski, D. Metzler, V. Murdock, and V. Plachouras. To swing or not to swing: Learning when (not) to advertise. In *Proceeding of the 17th ACM conference on Information and knowledge management*, pages 1003–1012. ACM, 2008.
- [Cla73] E. Clarke. Multipart pricing of public goods. In *Public Choice*, page 41:617631, 1973.
- [EOS05] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords, 2005.
- [GH03] A.V. Goldberg and J.D. Hartline. Envy-free auctions for digital goods. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 29–35. ACM, 2003.
- [GHK⁺06] A.V. Goldberg, J.D. Hartline, A.R. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [Gro71] T. Groves. Incentives in teams. In *Econometrica*,, pages 11:17–33, 1971.
- [GS10] A. Ghosh and A. Sayedi. Expressive auctions for externalities in online advertising. In *Proceedings of the 19th international conference on World wide web*, pages 371–380. ACM, 2010.

- [HIMM11] N. Haghpanah, N. Immorlica, VS Mirrokni, and K. Munagala. Optimal auctions with positive network externalities. In *Proc. of the 12th ACM Conference on Electronic Commerce (EC11)*, pages 11–20, 2011.
- [JMS96] P. Jehiel, B. Moldovanu, and E. Stacchetti. How (not) to sell nuclear weapons. *The American Economic Review*, pages 814–829, 1996.
- [JS11] K. Jerath and A. Sayedi. Exclusive display in sponsored search advertising. 2011.
- [LPL11] B. Lucier and R. Paes Leme. Gsp auctions with correlated types. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 71–80. ACM, 2011.
- [LT10] R.P. Leme and E. Tardos. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 735–744. IEEE, 2010.
- [Mut09] S. Muthukrishnan. Bidding on configurations in internet ad auctions. *Computing and Combinatorics*, pages 1–6, 2009.
- [PKK11] J. Pei, D. Klabjan, and F. Karaesmen. Auction Analysis for Sharable Goods with Externality. *Working Paper*, 2011.
- [SKN⁺10] S. Schroedl, A. Kesari, A. Nair, L. Neumeyer, and S. Rao. Generalized Utility in Web Search Advertising. 2010.
- [SS71] L.S. Shapley and M. Shubik. The assignment game i: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.
- [Vic61] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. In *Journal of Finance*, page 16:837, 1961.